Contrast Enhancement Techniques for Plant Leaf Disease Detection: A Review of CLAHE and Related Preprocessing Methods in CNN-based Diagnosis

R. Senthil
Research Scholar
Sangam University,
Bhilwara-311001, Rajasthan, India

R. Khatwal
Associate Professor
Sangam University,
Bhilwara-311001, Rajasthan, India

ABSTRACT
Low and uneven contrast in plant leaf images captured in the field continues to hinder the reliability of convolutional neural network (CNN)-based disease detection systems. This is because lesion boundaries, chlorotic margins, and early necrotic spots often exhibit only minor local intensity differences from healthy tissue. Contrast Limited Adaptive Histogram Equalization (CLAHE) has emerged as a widely adopted preprocessing step to address this issue and has been integrated with various architectures, from custom shallow CNNs to deep transfer-learning backbones such as ResNet50 and VGG16. This paper reviews the application of CLAHE and related contrast-enhancement techniques in the plant disease detection literature. It summarizes the algorithmic foundations of CLAHE and its predecessors, surveys recent studies that incorporate CLAHE into CNN and capsule-network pipelines for leaf disease classification, and compares the datasets, architectures, and reported results across these works. While CLAHE is consistently reported to improve feature visibility and subsequent classification performance, very few studies systematically vary or justify its two governing parameters—clip limit and tile grid size. Even fewer report the statistical significance of the resulting improvements. The review concludes by identifying specific research gaps, including the necessity for controlled parameter sweeps, baseline comparisons against alternative enhancement methods, and standardized statistical validation, all of which could be addressed in future controlled studies.

General Terms
Image Processing, Pattern Recognition, Machine Learning, Agricultural Informatics.

Keywords
Contrast Limited Adaptive Histogram Equalization, CLAHE, image enhancement, plant disease detection, convolutional neural network, ResNet50, transfer learning, agricultural image processing.

1. INTRODUCTION
Accurate and timely identification of plant leaf diseases is vital for minimizing crop losses and enhancing food security. Image-based deep learning methods have become the leading approach for automating this task [1][2]. Convolutional neural networks (CNNs), particularly transfer-learning architectures like ResNet50 [3] pretrained on ImageNet [4], consistently show strong classification performance on benchmark datasets such as PlantVillage [1][5]. However, performance on controlled laboratory datasets often does not translate effectively to field-captured images. In real-world settings, factors such as uneven illumination, shadows, and low local contrast between lesions and healthy tissue frequently obscure disease-relevant features [2][5]. Contrast enhancement is a natural preprocessing step for this problem. Contrast Limited Adaptive Histogram Equalization (CLAHE) [6] is a commonly employed technique in recent plant disease detection pipelines. It has been integrated with various architectures, ranging from lightweight custom CNNs to attention-augmented VGG16 [7] and capsule networks [8]. Despite its widespread use, the literature seldom treats CLAHE's hyperparameters—specifically, the clip limit (which restricts local contrast amplification) and the tile grid size (which determines the spatial resolution of adaptive equalization)—as variables requiring optimization or detailed reporting. This review synthesizes existing literature on CLAHE and related contrast-enhancement techniques utilized in plant leaf disease detection. It specifically examines how these studies configure, justify, and evaluate their preprocessing choices, thereby motivating and defining the scope for a subsequent controlled experimental study.

The remainder of this paper is organized as follows: Section 2 describes the review methodology. Section 3 reviews the algorithmic foundations of histogram-based contrast enhancement. Section 4 summarizes CNN and transfer-learning architectures employed for plant disease detection, including their most common associated datasets. Section 5 synthesizes the specific literature that combines CLAHE with these architectures, offering a quantitative synthesis and a comprehensive evaluation matrix across studies. Section 6 discusses explainability and evaluation practices within this literature. Section 7 identifies open research gaps, and Section 8 concludes the paper.

2. REVIEW METHODOLOGY
This review utilizes a narrative/scoping literature review approach rather than a fully quantified PRISMA-style systematic review. While it does not provide exhaustive database record counts at each screening stage, it employs an explicit and repeatable search and inclusion process, detailed below for transparency and to facilitate replication or extension.

2.1 Search Strategy and Sources
Table 1 outlines the sources and representative search terms used to identify literature integrating contrast enhancement with CNN-based plant disease detection. The search encompassed general web/scholar platforms, preprint servers, peer-reviewed journal databases, and applied/practitioner literature. Direct verification of the target dataset's provenance was also performed.

Table 1. Search sources and representative search terms

Source	Example Search Terms	Primary Use
Web / Google Scholar	"CLAHE plant disease", "contrast enhancement leaf"	General discovery
arXiv	"CLAHE CNN", "attention leaf disease"	Recent preprints
PLOS ONE / ScienceDirect	"CLAHE capsule network", "rice leaf review"	Peer-reviewed articles
IJRASET repository	"CLAHE custom CNN plant"	Applied / practitioner studies
Mendeley Data	"rice leaf diseases dataset"	Dataset provenance verification

2.2 Inclusion and Exclusion Criteria

Table 2 summarizes the criteria used to select studies for synthesis in Sections 4 and 5, including topical relevance, publication type, recency, and language.

Table 2. Study inclusion and exclusion criteria

Criterion	Included	Excluded
Topic	CNN / CapsNet leaf-disease classification with a stated contrast-enhancement step	Generic image enhancement unrelated to agriculture
Publication type	Peer-reviewed journals, established preprint servers	Unverifiable blogs, non-indexed sources
Recency	2015–2026, with pre-2015 foundational method papers kept	Superseded early enhancement techniques without modern relevance
Language	English	Non-English without English translation

3. CONTRAST ENHANCEMENT TECHNIQUES: BACKGROUND

3.1 Histogram Equalization

Global Histogram Equalization (HE) redistributes an image's intensity values to approximate a uniform cumulative distribution of pixel intensities. This process typically spreads out the most frequent intensity values, thereby increasing overall contrast. However, because HE computes a single transformation for the entire image, it can over-amplify noise in nearly uniform regions and under-enhance small, spatially localized features, such as early-stage lesions—structures critical for disease detection.

3.2 Adaptive Histogram Equalization and CLAHE

Adaptive Histogram Equalization (AHE) addresses the locality problem by computing separate histogram equalizations for small, overlapping, or tiled regions of the image and interpolating between them. This technique was originally proposed for medical image contrast enhancement [9]. A known drawback of AHE is its potential to over-amplify noise in homogeneous regions. To mitigate this effect, Zuiderveld [6] introduced Contrast Limited Adaptive Histogram Equalization (CLAHE). Before computing the cumulative distribution function for each tile, the local histogram in CLAHE is clipped at a predefined limit, and the clipped portion is redistributed across the histogram bins, as illustrated (see Figure 1). This clip limit directly controls the maximum local contrast amplification, while the tile grid size governs the spatial granularity of the adaptive equalization. By operating independently on each tile prior to bilinear interpolation for boundary smoothing, CLAHE enhances local contrast—such as the boundary between a lesion and surrounding healthy tissue—without the global over-amplification associated with plain HE.

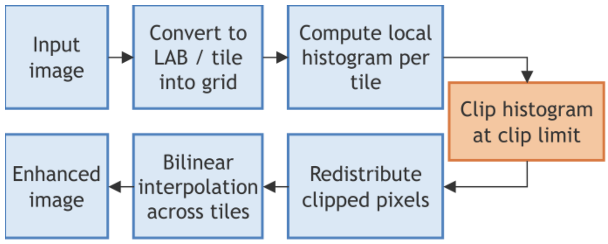

Figure 1: Block diagram of the CLAHE algorithm pipeline. The clip-limit step bounds local contrast amplification per tile.

Note: The clip limit constrains local contrast amplification within each image tile, whereas the tile grid size determines the spatial resolution of the adaptive histogram equalization process.

3.3 Other Enhancement Approaches and Technique Comparison

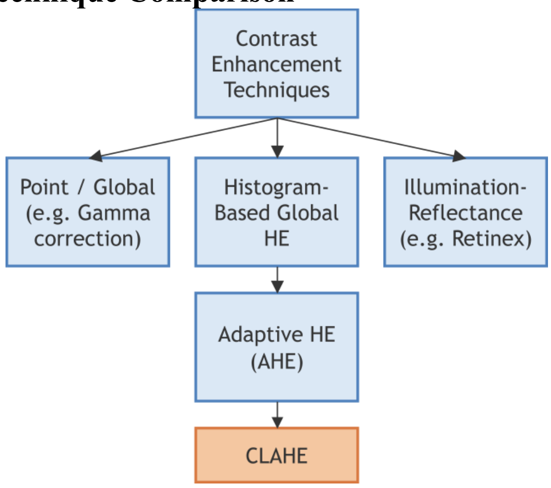

Figure 2: Taxonomy of contrast-enhancement techniques referenced in the surveyed literature, tracing CLAHE's lineage from global HE through AHE.

Beyond HE and CLAHE, the broader image enhancement literature offers other techniques such as gamma correction and Retinex-based methods.Gamma correction adjusts image brightness or darkness using a power-law intensity transformation but does not account for local image structures. Retinex-based methods, conversely, aim to separate illumination from reflectance to standardize lighting variations. (see Figure 2) it integrates these techniques into a unified taxonomy, while Table 3 compares their locality, governing parameters, and significance in the reviewed literature.within the plant disease detection literature, CLAHE is the most frequently reported enhancement technique. It is typically combined with complementary feature-domain techniques, such as color-difference histograms [8], rather than being directly compared against alternative enhancement methods.

Table 3. Comparison of contrast-enhancement technique properties

Technique	Locality	Key Parameters	Role in Reviewed Literature
Gamma correction	Global (point-wise)	Gamma exponent	Rarely primary enhancement in surveyed studies
Global HE	Whole image	None	Baseline comparator; can over-amplify noise
Adaptive HE (AHE)	Local tiles, uncapped	Tile size	Predecessor to CLAHE [9].
CLAHE	Local tiles, capped	Clip limit, tile grid	Dominant choice across all studies [6][7][8]
Retinex-based	Illumination / reflectance	Scale, weighting	Not used in reviewed CLAHE studies

4. DEEP LEARNING FOR PLANT DISEASE DETECTION

4.1 CNN and Transfer Learning Approaches

Mohanty et al. [1] demonstrated the technical feasibility of large-scale, CNN-based plant disease classification. They achieved over 99% accuracy in controlled-condition evaluations by training models on 54,306 images from the PlantVillage dataset [5], which encompasses 26 diseases across 14 crop species. Ferentinos [2] further explored this by employing a broader array of CNN architectures and a combination of laboratory and field images. This research identified a significant decrease in accuracy when models trained on laboratory images were evaluated with field-captured images, directly attributing this to differences in illumination and contrast between the two capture environments. Too et al. [10] systematically compared fine-tuning strategies for VGG16, Inception V4, ResNet (with 50, 101, and 152 layers), and DenseNet121 on the PlantVillage benchmark. They concluded that deeper, more parameter-efficient architectures like DenseNet121 generalized better with less overfitting than shallower networks. ResNet-family models also remained competitive and were considerably more cost-effective to fine-tune than to train from scratch.

4.2 ResNet50 as a Backbone for Leaf Disease Classification

ResNet50 [3] introduced residual (skip) connections, which enable stable training of very deep networks by facilitating gradient propagation and preventing vanishing gradients. Due to its availability as an ImageNet-pretrained [11] backbone, ResNet50 is frequently selected for transfer learning in plant disease classification. The standard fine-tuning procedure involves replacing and retraining the final classification layer, with the option to unfreeze deeper convolutional blocks [12]. Table 4 compares ResNet50 to other CNN and capsule-network backbones discussed in this paper, detailing their depth, parameter count, and core mechanisms.

Table 4. CNN and capsule-network backbone architectures referenced across the reviewed literature

Architecture	Depth	Params	Key Mechanism
VGG16 [13] (2014)	16	~138M	Stacked 3x3 convolutions
ResNet50 [3] (2016)	50	~25.6M	Residual connections
DenseNet121 [14] (2017)	121	~8M	Dense connectivity
Inception-v4 (2016–17)	~76–164	~35–55M	Multi-scale convolutions
CapsNet-based [8]	Shallow	~2–5M	Capsule routing

4.3 Benchmark Datasets

PlantVillage [5] is the most widely used benchmark for classifying plant diseases, containing over 54,000 images across 38 crop-disease categories, all captured under controlled conditions. Simhadri et al. [15] document an increasing number of rice leaf disease datasets and observe that most studies since 2017 employ transfer learning, ensemble learning, or hybrid architectures. They also highlight that variations in image quality and preprocessing methods consistently contribute to performance differences across studies. Table 5 give the specific datasets referenced in this review, including the exact per-subset image counts reported by Singh et al. [7]. The Rice Leaf Diseases dataset [16], the target dataset for this review's experimental study, contains 4,684 images distributed across three classes: Bacterial Blight, Brown Spot, and Leaf Smut. Like other rice leaf collections acquired in the field, this dataset exhibits uneven illumination and contrast variability, underscoring the need for CLAHE-based preprocessing.

Table 5. Benchmark datasets referenced in this review

Dataset	Cls.	Images	Cond.	Ref.
PlantVillage [1][5]	38	54,306	Ctrl.	[1][7][8] [10]
Embrapa (via [7])	93	46,376	Field	[7]
Maize (via [7])	4	3,852	Field	[7]
Apple (via [7])	4	3,644	Field	[7]
Rice (via [7])	10	9,000	Field	[7]
Rice Leaf Dis. [16]	3	4,684	Field	Target

5. CLAHE IN PLANT DISEASE DETECTION: A LITERATURE SYNTHESIS

5.1 CLAHE as a Preprocessing Step for CNN Classifiers

Recent research commonly integrates CLAHE as a standard preprocessing step before training Convolutional Neural Network (CNN) classifiers, rather than evaluating it as an experimental variable. For example, a study in IJRASET utilized CLAHE to process the PlantVillage dataset before training a custom CNN. This approach yielded a training accuracy of 94.64%, a validation accuracy of 92.95%, and a loss of 0.1793. The authors attributed the improved feature visibility directly to CLAHE's contrast enhancement, although they did not conduct an ablation study to compare it against a non-CLAHE baseline. Similarly, Singh et al. [7] made CLAHE the initial stage in their five-dataset preprocessing pipeline (detailed in Table 5), which included a rice subset. This was followed by an attention-augmented VGG16 classifier, achieving accuracies up to 98.87% on the rice dataset. However, this study also did not isolate CLAHE's specific contribution from that of the attention mechanism. (see Figure 3) presents a generalized reference architecture, synthesized from the pipelines described in these reviewed studies.

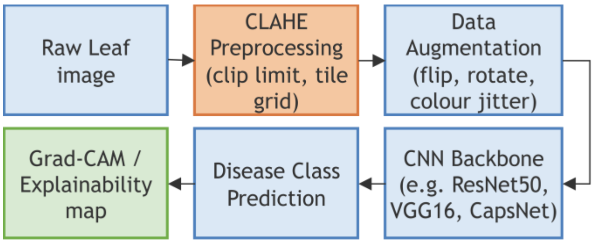

Figure 3: General pipeline of CLAHE-based CNN plant leaf disease detection systems, synthesized from the studies reviewed in Section 5. Shaded stages are the least rigorously parameterized in current practice

5.2 CLAHE Combined with Complementary Feature Extraction

Okyere-Gyamfi et al. [8] introduced CCFM-CapsNet, a capsule network that integrated CLAHE with a color-difference histogram (CDH) feature extractor. CLAHE was incorporated for its capacity to standardize local contrast and mitigate sensitivity to lighting variations, particularly when foreground and background exhibit similar overall brightness. The model was assessed using ten crop-disease subsets from PlantVillage, including rice, achieving validation accuracies ranging from 93% to 100%, contingent on the specific subset. Consistent with previously discussed CNN-based studies, CLAHE's clip limit and tile grid size were fixed rather than optimized. Moreover, their individual impact on performance was not disaggregated from the contribution of the CDH feature merging.

5.3 CLAHE Parameter Choices Across the Literature

Table 6 summarizes representative studies that integrate CLAHE preprocessing with CNN or capsule-network classifiers for leaf disease detection, outlining their reported architecture and primary results. (Dataset specifics for each study are presented in Table 5). Table 7 then meticulously details how each study documented CLAHE's hyperparameters. A consistent pattern emerges: CLAHE's clip limit and tile grid size are typically treated as predetermined implementation details, frequently adopted from default library values (e.g., OpenCV's default clip limit of 2.0 and an 8x8 tile grid), rather than being optimized hyperparameters tailored to the specific dataset and downstream architecture. None of the reviewed studies reported a systematic grid search across both parameters concurrently, nor did they offer a baseline comparison against a no-enhancement or histogram-equalization control, trained and evaluated under otherwise identical conditions.

Table 6. Representative studies combining CLAHE with CNN/capsule-network classifiers for leaf disease detection

Study	Arch.	Enhance.	Result
Singh [7]	CBAM-VGG16	CLAHE + norm.	up to 98.87%
Okyere-Gyamfi [8]	CCFM-CapsNet	CLAHE + CDH	93.5–100%
IJRASET study	Shallow CNN	CLAHE	94.64 / 92.95%
This review	ResNet50	CLAHE sweep	open gap

Table 7. CLAHE hyperparameter reporting across the surveyed studies

Study	Clip Limit?	Tile Grid?	Param. Sweep?
Singh et al. [7]	No	No	No
Okyere-Gyamfi et al. [8]	No	No	No
IJRASET custom-CNN	No	No	No
Library default (OpenCV)	2.0 (default)	8x8 (default)	N/A
This review (target)	1–4 sweep	4x4 / 8x8 / 16x16	Yes

5.4 Reported Accuracy Across Enhancement Approaches

Table 8 presents the primary accuracy figures reported in Sections 5.1–5.2, comparing them with similar figures from studies in Section 4.1 that did not incorporate an explicit enhancement step. While this is not a controlled comparison—due to variations in underlying architectures, datasets, and evaluation protocols among studies—it indicates that reported accuracies in both groups generally range from the low to upper 90s.

Table 8. Reported Headline Accuracy by Enhancement Approach Across Surveyed Studies

Enhancement Approach	Representative Study	Reported Accuracy
No stated enhancement	Mohanty et al. [1] (PlantVillage, controlled)	>99% (controlled only)
No stated enhancement	Too et al. [10] (PlantVillage, fine-tuned CNNs)	High-90s%, varies by backbone
CLAHE + attention CNN	Singh et al. [7] (rice subset)	Up to 98.87%
CLAHE + CapsNet / CDH	Okyere-Gyamfi et al. [8] (subset-dependent)	93.5–100%
CLAHE + custom CNN	IJRASET study (PlantVillage)	92.95% (val.)

5.5 Quantitative Synthesis Across Studies

(see Figure 4) visually consolidates the accuracy figures from Table 8, offering a clearer comparison of headline results across studies than the table alone. Due to variations in datasets, backbone architectures, class counts, and evaluation protocols among the underlying studies, this comparison describes the current reporting landscape rather than providing evidence of CLAHE's causal effect on accuracy.

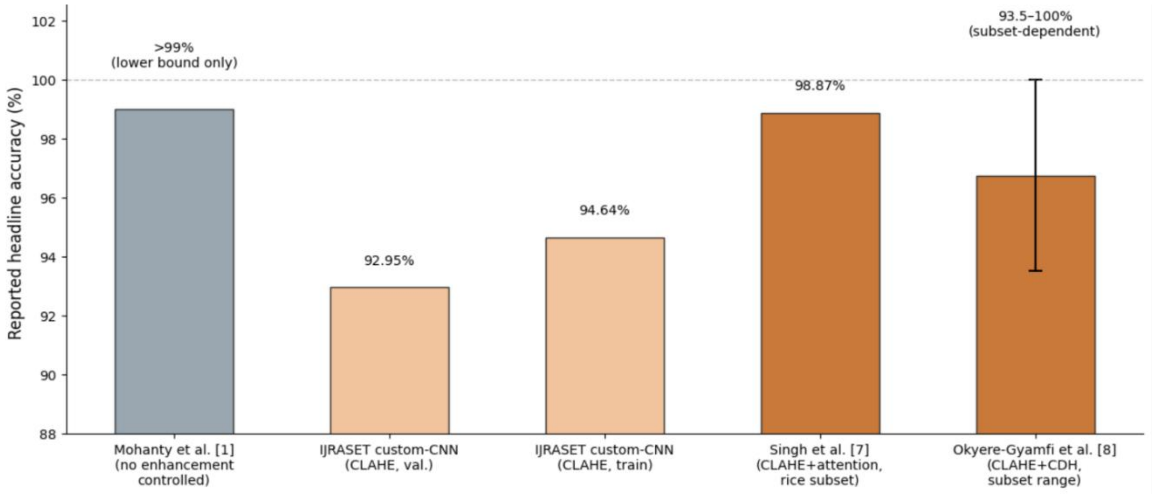

Figure 4: Reported headline accuracy by study, consolidated from Table 8. Bars show point estimates as originally reported; the error bar for Okyere-Gyamfi et al. [8] shows the published subset range rather than a statistical confidence interval.

(see Figure 4) reveals three descriptive patterns. First, all headline accuracies cluster within a narrow range of approximately 93% to 100%, irrespective of whether the pipeline included a stated contrast-enhancement step. This narrow spread reinforces the review's consistent interpretation: reported accuracy alone is insufficient to isolate CLAHE's specific contribution; a controlled, single-variable comparison is needed, not a between-study one. Second, the two studies that combine CLAHE with an additional mechanism—attention modules in Singh et al. [7] and color-difference histogram feature merging in Okyere-Gyamfi et al. [8]—report accuracies at or above the upper end of the observed range. However, neither study includes an ablation to isolate CLAHE's marginal contribution from that of the co-occurring mechanism. Third, Okyere-Gyamfi et al. [8] exhibits the widest reported spread (93.5–100%), which reflects variation across the ten crop-disease subsets evaluated, not variation attributable to contrast enhancement itself. This subset-level heterogeneity cannot be resolved from the published headline figures alone. Since none of the reviewed studies report standard deviations, confidence intervals, or repeated-run statistics alongside their headline figures, (see Figure 4) intentionally avoids implying a statistical comparison between bars. The range shown for Okyere-Gyamfi et al. [8] represents the published subset range, not a confidence interval, and the marker for Mohanty et al. [1] indicates a reported lower bound (">99%") rather than an exact value. This lack of dispersion measures is a central finding of this review and is formally addressed as a research gap in Section 7.

5.6 Comprehensive Evaluation Matrix

Table 9 consolidates the criteria from Tables 4–8 into a single, comprehensive evaluation matrix. This matrix allows for a unified inspection of each surveyed study's dataset characteristics, backbone architecture, CLAHE hyperparameter reporting, methodological controls, and evaluation practices, rather than examining them across several tables. (see Figure 5) illustrates the hyperparameter-reporting dimension of this matrix, specifically clip limit, tile grid, and parameter-sweep reporting, thereby extending the tabular presentation in Table 7.

Table 9 and (see Figure 5) collectively show that no surveyed empirical study integrates all four of the following properties: (i) documented CLAHE hyperparameters, (ii) a stated non-CLAHE baseline trained under matched conditions, (iii) any form of statistical testing, and (iv) an explainability check specifically targeting the contrast-enhancement step rather than the classifier as a whole. The "library-default" row (OpenCV) functions as a reference point, not an empirical study, illustrating that while default parameter values are well-documented in software, this documentation does not inherently justify their application for a specific dataset and task. The final row outlines the planned methodological approach of the target study introduced in Section 7, which is designed to concurrently satisfy all four criteria. Items marked "planned" represent proposed future work and are not yet supported by completed experiment.

Table 9. Comprehensive evaluation matrix across surveyed studies

Study	Dataset	Backbone	Enhancement	Clip Limit	Tile Grid	Param. Sweep	Non-CLAHE Baseline	Stat. Testing	Explainability
Singh et al. [7]	5 field dataset (Table 5)	CBAM-VGG16	CLAHE + norm.	No	No	No	No	No	Yes (Grad-CAM family)
Okyere-Gyamfi et al. [8]	PlantVillage, 10 subsets	CCFM-CapsNet	CLAHE + CDH	No	No	No	No	No	No
IJRASET custom-CNN	PlantVillage	Custom shallow CNN	CLAHE	No	No	No	No	No	No
Library default (OpenCV)	N/A	N/A	CLAHE (default)	Yes (2.0)	Yes (8x8)	N/A	N/A	N/A	N/A

This review (target)	Rice Leaf [16]	ResNet50	CLAHE sweep vs. HE vs. none	Yes (1–4)	Yes (4x4/8x8/16x16)	Yes (planned)	Planned	Planned	Planned

Figure 5: CLAHE hyperparameter-reporting practice across surveyed studies, visualized from Table 7 and Table 9.

6. EXPLAINABILITY AND EVALUATION PRACTICES

6.1 Grad-CAM and Related Visualizations

Gradient-weighted Class Activation Mapping (Grad-CAM) [17] generates coarse localization maps that identify image regions most influential in a Convolutional Neural Network's (CNN) prediction. This is achieved by weighting the final convolutional layer's feature maps with the gradient of the target class score. Singh et al. [7] employed Grad-CAM, Grad-CAM++, and layer-wise relevance propagation to visualize the leaf regions emphasized by their CBAM-VGG16 model. They found that the resulting maps focused on lesion margins and discolored patches, rather than background or vein structures. While this offers a valuable qualitative indicator of model validity, it does not substitute for quantitative ablation of the preprocessing pipeline itself, as shown in the shaded stage (see Figure 3). In the reviewed CLAHE-specific literature, explainability visualization primarily validates the final classifier, rather than directly diagnosing which contrast-enhancement configuration most effectively improves feature saliency.

6.2 Data Augmentation Practices

Data augmentation, encompassing techniques such as geometric transformations and color jitter, is a ubiquitous practice in this field. It complements contrast enhancement by mitigating the challenges of relatively small plant disease datasets compared to the complexity of modern CNN backbones [18]. Across surveyed CLAHE-based studies, flipping, rotation, and zooming are consistently employed, typically after contrast enhancement in the preprocessing pipeline (see Figure 3). However, because augmentation and enhancement are usually varied together rather than independently, it is difficult to isolate their individual contributions to reported accuracy gains from published results alone.

7. OPEN CHALLENGES AND RESEARCH GAPS

Three recurring gaps are evident across the reviewed literature

(directly visible in Table 7 and summarized as a rigor scorecard in Table 10):

First, CLAHE's governing parameters—clip limit and tile grid size—are consistently fixed at default or lightly tuned values rather than optimized through a documented search. This raises questions about the sensitivity of downstream classification performance to these choices for a given dataset and architecture.Second, few studies report a controlled baseline comparison that isolates CLAHE's contribution from other pipeline components such as attention modules, augmentation, or alternative feature extractors. This makes it difficult to attribute reported accuracy gains specifically to contrast enhancement.Third, statistical validation—including multi-seed repetition, confidence intervals, or significance testing against a non-CLAHE baseline—is largely absent from the CLAHE-in-plant-disease literature reviewed here, in contrast to more mature areas of machine learning evaluation practice. Addressing these three gaps through a controlled sweep of clip-limit and tile-grid-size parameters, evaluated against matched no-enhancement and histogram-equalization baselines with multi-seed statistical testing, represents a concrete and tractable direction for a follow-up experimental study on the Rice Leaf Diseases dataset [16] discussed in Section 4.3.

Table 10. Methodological rigor scorecard for the surveyed CLAHE-based studies

Study	CLAHE Param. Search	Non-CLAHE Baseline	Stat. Testing
Singh et al. [7]	No	No	No

Okyere-Gyamfi et al. [8]	No	No	No
IJRASET custom-CNN	No	No	No
This review (target)	Yes	Yes	Yes

8. CONCLUSION

CLAHE has emerged as a standard preprocessing technique in contemporary plant leaf disease detection pipelines, consistently enhancing classification accuracy across various architectures, including Convolutional Neural Networks, attention-augmented models, and capsule networks. Nevertheless, the existing literature frequently treats CLAHE's parameters as immutable implementation details rather than variables necessitating rigorous optimization and evaluation. Moreover, CLAHE's specific impact is seldom isolated from other pipeline components through controlled baselines or robust statistical testing. A systematic investigation into the optimal configuration of clip-limit and tile-grid-size, benchmarked against carefully matched no-enhancement and histogram equalization baselines, and validated through multi-seed statistical analysis on a ResNet50 backbone, would directly address this critical research gap and represents the logical progression indicated by this review.

9. ACKNOWLEDGMENTS

The authors extend their gratitude to Sangam University, particularly the Faculty of Computer Science and Engineering, for their support. This work contributes to the first author's doctoral training and partially fulfills the requirements for thesis preparation.

10. REFERENCES

[1] S. P. Mohanty, D. P. Hughes, and M. Salathé, “Using Deep Learning for Image-Based Plant Disease Detection,” Front. Plant Sci., vol. 7, p. 1419, Sep. 2016, doi: 10.3389/fpls.2016.01419.

[2] K. P. Ferentinos, “Deep learning models for plant disease detection and diagnosis,” Comput. Electron. Agric., vol. 145, pp. 311–318, Feb. 2018, doi: 10.1016/j.compag.2018.01.009.

[3] K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA: IEEE, Jun. 2016, pp. 770–778. doi: 10.1109/CVPR.2016.90.

[4] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” Commun. ACM, vol. 60, no. 6, pp. 84–90, May 2017, doi: 10.1145/3065386.

[5] David. P. Hughes and M. Salathe, “An open access repository of images on plant health to enable the development of mobile disease diagnostics,” 2015, arXiv. doi: 10.48550/ARXIV.1511.08060.

[6] K. Zuiderveld, “Contrast Limited Adaptive Histogram Equalization,” in Graphics Gems, Elsevier, 1994, pp. 474–485. doi: 10.1016/B978-0-12-336156-1.50061-6.

[7] B. Singh, R. P. Sharma, and S. Dey, “Interpretable Plant Leaf Disease Detection Using Attention-Enhanced CNN,” 2025, arXiv. doi: 10.48550/ARXIV.2512.17864.

[8] S. Okyere-Gyamfi, M. Asante, K. O. Peasah, Y. M. Missah, and V. Akoto-Adjepong, “Contrast limited adaptive histogram equalization (CLAHE) and colour difference histogram (CDH) feature merging capsule network (CCFMCapsNet) for complex image recognition,” PLOS One, vol. 20, no. 10, p. e0335393, Oct. 2025, doi: 10.1371/journal.pone.0335393.

[9] S. M. Pizer et al., “Adaptive histogram equalization and its variations,” Comput. Vis. Graph. Image Process., vol. 39, no. 3, pp. 355–368, Sep. 1987, doi: 10.1016/S0734-189X(87)80186-X.

[10] E. C. Too, L. Yujian, S. Njuki, and L. Yingchun, “A comparative study of fine-tuning deep learning models for plant disease identification,” Comput. Electron. Agric., vol. 161, pp. 272–279, Jun. 2019, doi: 10.1016/j.compag.2018.03.032.

[11] J. Deng, W. Dong, R. Socher, L.-J. Li, Kai Li, and Li Fei-Fei, “ImageNet: A large-scale hierarchical image database,” in 2009 IEEE Conference on Computer Vision and Pattern Recognition, Miami, FL: IEEE, Jun. 2009, pp. 248–255. doi: 10.1109/CVPR.2009.5206848.

[12] S. J. Pan and Q. Yang, “A Survey on Transfer Learning,” IEEE Trans. Knowl. Data Eng., vol. 22, no. 10, pp. 1345–1359, Oct. 2010, doi: 10.1109/TKDE.2009.191.

[13] K. Simonyan, A. Vedaldi, and A. Zisserman, “Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps,” 2013, arXiv. doi: 10.48550/ARXIV.1312.6034.

[14] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, “Densely Connected Convolutional Networks,” in 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI: IEEE, Jul. 2017, pp. 2261–2269. doi: 10.1109/CVPR.2017.243.

[15] C. G. Simhadri, H. K. Kondaveeti, V. K. Vatsavayi, A. Mitra, and P. Ananthachari, “Deep learning for rice leaf disease detection: A systematic literature review on emerging trends, methodologies and techniques,” Inf. Process. Agric., vol. 12, no. 2, pp. 151–168, Jun. 2025, doi: 10.1016/j.inpa.2024.04.006.

[16] L. P. Lourdu Antony, “Rice Leaf Diseases Dataset.” Mendeley Data, Oct. 19, 2023. doi: 10.17632/DWTN3C6W6P.1.

[17] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization,” Int. J. Comput. Vis., vol. 128, no. 2, pp. 336–359, Feb. 2020, doi: 10.1007/s11263-019-01228-7.

[18] C. Shorten and T. M. Khoshgoftaar, “A survey on Image Data Augmentation for Deep Learning,” J. Big Data, vol. 6, no. 1, p. 60, Dec. 2019, doi: 10.1186/s40537-019-0197-0.